arXiv:1411.0924v2 [stat.AP] 18 Apr 2016

# An adaptive spatio-temporal smoothing model for estimating trends and step changes in disease risk

Alastair Rushworth[a], Duncan Lee[b] and Christophe Sarran[c]

[a] *Department of Mathematics and Statistics, University of Strathclyde, UK.*

[b] *School of Mathematics and Statistics, University of Glasgow, UK.*

[c] *UK Met Office, Exeter, UK.*

**Summary**. Statistical models used to estimate the spatio-temporal pattern in disease risk from areal unit data represent the risk surface for each time period with known covariates and a set of spatially smooth random effects. The latter act as a proxy for unmeasured spatial confounding, whose spatial structure is often characterised by a spatially smooth evolution between some pairs of adjacent areal units while other pairs exhibit large step changes. This spatial heterogeneity is not consistent with existing global smoothing models, in which partial correlation exists between all pairs of adjacent spatial random effects. Therefore we propose a novel space-time disease model with an adaptive spatial smoothing specification that can identify step changes. The model is motivated by a new study of respiratory and circulatory disease risk across the set of Local Authorities in England, and is rigorously tested by simulation to assess its efficacy. Results from the England study show that the two diseases have similar spatial patterns in risk, and exhibit a number of common step changes in the unmeasured component of risk between neighbouring local authorities.

**Keywords: Adaptive smoothing; Gaussian Markov random fields; Spatio-temporal disease mapping; Step change detection.**

## 1. Introduction

Disease risk exhibits spatio-temporal variation due to many factors, including changing levels of environmental exposures and risk-inducing behaviours such as smok-

†*Email address for correspondence: alastair.rushworth@strath.ac.uk*

ing. Disease risk data are typically obtained as population level summaries for administrative geographical units, such as local authorities, and the spatial pattern in risk is presented via a choropleth map. Such maps enable public health professionals to quantify the spatial pattern in disease risk, allowing financial resources and public health interventions to be targeted at high-risk areas. Disease maps are routinely published by health agencies worldwide, such as the cancer e-Atlas (`http://www.ncin.org.uk/cancer_information_tools/eatlas/`) by Public Health England and the weekly influenza maps (`http://www.cdc.gov/flu/weekly/usmap.htm`) produced by the Centres for Disease Control and Prevention in the USA. These maps also allow the scale of health inequalities between rich and poor communities and their underlying drivers to be quantified. For example, a 2014 report by the Office for National Statistics (2014) estimates that UK average healthy life expectancy differs by nineteen years between communities with the highest and the lowest levels of deprivation. Such large inequalities exacerbate socioeconomic divisions in society, and health costs may be increased due to higher disease prevalence in the most disadvantaged regions.

Disease maps presented by health agencies display raw disease rates, which do not allow inferential statements to be made such as calculating risk-exceedence probabilities (Richardson et al., 2004), or evaluating the significance of temporal changes. A range of statistical models have been developed for disease data, which represent the risk surface using known covariates and a set of random effects. The latter act as a proxy for unmeasured spatial confounding, and are typically modelled by a Gaussian Markov Random field (GMRF, Rue and Held, 2005) prior in a hierarchical framework. GMRF priors have been extended to incorporate spatio-temporal structure, and prominent examples include Bernardinelli et al. (1995), Knorr-Held (2000), MacNab and Dean (2001) and Ugarte et al. (2010).

GMRF models assume the random effects are globally spatially smooth, in the sense that a single parameter governs the spatial autocorrelation in disease risk between all pairs of geographically adjacent units. In practice however, this residual or un-

explained spatial structure is often characterised by a spatially smooth evolution between some pairs of adjacent units, while other pairs exhibit large step changes. The identification of such step changes in the unexplained component of risk is known as *Wombling*, following the seminal article by Womble (1951), and can provide a number of epidemiological insights. Firstly, it allows the identification of the geographical extent of clusters of areal units that exhibit elevated unexplained risks, which enables health resources and public health interventions to be targeted appropriately. Secondly, it provides detailed insight into the spatial structure of the unmeasured confounding, allowing the identification of unknown etiological factors that contribute to disease risk. Furthermore, smoothing models that ignore local structure may result in oversmoothing in regions where strong disparities exist and undersmoothing elsewhere, leading to biased estimation of disease risk. A range of spatially adaptive smoothing priors have been proposed to address these limitations for purely spatial data, including Green and Richardson (2002), Lu et al. (2007), Lawson et al. (2012), Lee and Mitchell (2013), Wakefield and Kim (2013) and Lee et al. (2014).

Few spatially adaptive smoothing models have been developed for spatio-temporal disease data, with an exception being Lee and Mitchell (2014) who propose an iterative fitting algorithm using integrated nested Laplace approximations. While temporal replication is likely to improve the estimation in such highly complex models, the increased numbers of data points and parameters results in much increased computational complexity. Therefore, the contribution of this paper is the development of a computationally feasible spatially-adaptive GMRF model for spatio-temporal disease data, which can be viewed as both an adaptive smoother and a model for the detection of step changes in unexplained risk. The model builds on the purely spatial approach of Ma et al. (2010), and avoids making simplifying assumptions about the step-change structure as Lee et al. (2014) do. Additionally, our model is freely available via the `R` package `CARBayesST` (Lee et al., 2015). The methodological development is motivated by a new study of respiratory and circulatory disease in England, UK, which according to the World Health Organisation (WHO) are two of the largest

causes of death worldwide (`www.who.int/mediacentre/factsheets/fs310/en/`). This study is presented in Section 2, while in Section 3 the literature on spatio-temporal disease mapping is reviewed. Section 4 proposes a new space-time GMRF model for adaptive smoothing, which is comprehensively tested by simulation in Section 5. In Section 6 the proposed model is applied to the motivating application, while the paper concludes in Section 7.

## 2. Motivating case study

Our methodological development is motivated by a new study of circulatory and respiratory disease risk in England between 2001 and 2010, which have International Classification of Disease tenth revision codes I00-I99 and J00-J99 respectively. Hospital admissions records from the Health and Social Care Information Centre (HSCIC) were analysed at the UK Met Office to provide yearly counts of emergency admissions by local and unitary authority (LUA). The resulting data $\{Y_{ij}\}$ are counts of hospital admissions for $i = 1, \ldots, N$ ($N = 323$) LUA in England in year $j = 1, \ldots, T$ ($T = 10$), and range between 6 - 1030 (circulatory) and 0 - 2485 (respiratory) respectively. The expected numbers of hospital admissions $\{E_{ij}\}$ were calculated for each year and LUA to adjust for their differing population sizes and demographic structures using indirect standardisation. Specifically, $E_{ij} = \sum_{r=1}^{q} N_{ijr} p_r$, where $N_{ijr}$ is the population size in LUA $i$, year $j$ and strata $r$ (e.g. males 0-5, etc), while $p_r$ is the England-wide risk of disease in strata $r$. The expected counts range between 144.8 - 564.1 (circulatory) and 115.2 - 675.9 (respiratory) respectively.

The Standardised Incidence Ratio, $\text{SIR}_{ij}$=$Y_{ij}/E_{ij}$ is an exploratory (noisy) measure of disease risk, and a value of 1.2 corresponds to a 20% increased risk of disease relative to $E_{ij}$. The mean SIR over all years is displayed in the left column of Figure 1, and similar spatial patterns are evident between the two diseases with a Pearson's correlation coefficient of 0.9356. Within each map, the risk levels are spatially smooth across much of England, although a number of step changes are

visible, including around the cities of Birmingham and Manchester. There are a number of potential drivers of this spatial variation in disease risk, including socio-economic deprivation (poverty), air pollution and the differences between urban and rural areas. We measure poverty by the percentage of the working age population who are in receipt of Job Seekers Allowance (JSA) obtained from the HSCIC, while we obtain modelled particulate matter concentrations ($PM_{10}$) from the Department for Environment Food and Rural Affairs on a 1km grid which are then averaged to the LUA scale. Finally, the urban or rural nature of each LUA is measured by the proportion of middle layer super-output areas (MSOAs) classified as urban within each LUA. The (mean over year) residuals from a simple Poisson log-linear model with these covariates are displayed in the right column of Figure 1. The residual unexplained spatial variation in disease risk is autocorrelated for both diseases, with significant (at the 5% level) Moran's I statistics ranging between 0.204 and 0.328 across the years. However, Figure 1 also highlights that these unexplained spatial structures exhibit step changes, which are not compatible with a global spatial smoothing model.

The aims of modelling these data are twofold. First, we want to produce the best estimate of the spatio-temporal patterns in circulatory and respiratory disease risks, so that the extent of the health inequalities in these two diseases can be identified. Second, we wish to estimate the locations of the step changes in the unexplained risk surface (Wombling), so that the geographical extent of clusters of excessively high unexplained risks regions can be identified and investigated for possible causes. To achieve these goals we propose an adaptive smoothing model in Section 4, but first present a review of the literature in Section 3. In this paper we take a univariate approach and model each disease separately, as is standard practice in the space-time disease mapping literature. An alternative would be a bivariate spatio-temporal model, and the choice between a univariate or bivariate approach depends on the questions one wishes to address. Here our interest is in developing one of the first spatially adaptive spatio-temporal smoothing models for disease risk, and comparing its efficacy to global smoothing alternatives. Such alternatives do not exist in a

bivariate context, and thus we would be unable to assess the improvement our adaptive smoothing model provides in this setting. Furthermore, our interest is in estimating the presence, and similarity between diseases, of step changes in the unexplained component of disease risk, and a bivariate model would borrow strength across diseases and thus bias the results towards finding step changes common to both diseases.

## 3. Spatio-temporal disease mapping

The observed and expected disease counts for LUA $i$ and year $j$ are denoted by $(Y_{ij}, E_{ij})$ respectively, and the following Poisson log-linear model is commonly specified for these data:

$$\begin{aligned} Y_{ij}|E_{ij}, R_{ij} &\sim \text{Poisson}(E_{ij}R_{ij}) \quad i = 1, \ldots, N, \quad j = 1, \ldots, T, \\ \ln(R_{ij}) &= \mathbf{x}_{ij}^{\top}\boldsymbol{\beta} + \phi_{ij}, \\ \beta_r &\sim \text{N}(0, 10000) \quad r = 1, \ldots, p. \end{aligned} \tag{1}$$

Disease risk is represented by $R_{ij}$, which is on the same scale as the SIR. The log-risk $\ln(R_{ij})$ is modelled by a vector of $p$ known covariates $\mathbf{x}_{ij} = (x_{ij1}, \ldots, x_{ijp})$ with parameters $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)$, and a spatio-temporal random effect $\phi_{ij}$. GMRF priors are commonly used to induce spatial smoothness between the random effects, via a binary $N \times N$ adjacency matrix $\mathbf{W}$. Element $w_{ik} = 1$ if areas $i$ and $k$ share a common border (denoted $i \sim k$) and $w_{ik} = 0$ otherwise (denoted $i \nsim k$), while $w_{ii} = 0$ for all $i$. Numerous GMRF priors have been developed for purely spatial random effects $(\phi_1, \ldots, \phi_N)$, and the proposal by Leroux et al. (1999) has an attractive full conditional decomposition given by

$$\phi_i|\boldsymbol{\phi}_{-i}, \rho, \tau^2, \mathbf{W} \sim \text{N}\left(\frac{\rho\sum_{k=1}^{N} w_{ik}\phi_k}{\rho\sum_{k=1}^{N} w_{ik} + 1 - \rho}, \frac{\tau^2}{\rho\sum_{k=1}^{N} w_{ik} + 1 - \rho}\right), \tag{2}$$

where $\boldsymbol{\phi}_{-i} = (\phi_1, \ldots, \phi_{i-1}, \phi_{i+1}, \ldots, \phi_N)$. The conditional expectation of $\phi_i$ is a

**Fig. 1.** Left: Mean standardised incidence ratios (SIR) for circulatory (top left) and respiratory (bottom left) hospital admissions across English local authorities between 2001 and 2010. Right: Mean standardised Pearson residuals from a Poisson log-linear model with JSA, urbanicity and $PM_{10}$ as covariates for circulatory (top right) and respiratory admissions (bottom right). In all cases, the locations of four major English cities are indicated with arrows.

weighted average of adjacent $\phi_k$ (as specified by $\mathbf{W}$), which induces smoothness across the surface. Spatial smoothing is controlled by $\rho \in [0,1]$, where $\rho = 1$ corresponds to the intrinsic autoregressive model (Besag, 1974), while $\rho = 0$ corresponds to independence. The resulting joint distribution is given by $(\phi_1, \ldots, \phi_N) \sim \mathrm{N}(\mathbf{0}, \tau^2\mathbf{Q}(\rho, \mathbf{W})^{-1})$, where the precision matrix is given by $\mathbf{Q}(\rho, \mathbf{W}) = \rho[\mathrm{diag}(\mathbf{W1}) - \mathbf{W}] + (1-\rho)\mathbf{I}$, where $\mathbf{1}$ is an $N \times 1$ vector of ones and $\mathbf{I}$ is the $N \times N$ identity matrix.

Many spatio-temporal GMRF models have been developed in the disease mapping literature, with the first being Bernardinelli et al. (1995) who model $R_{ij}$ with linear time-trends that have spatially varying slopes and intercepts. In contrast, Knorr-Held (2000) introduced a decomposition of $R_{ij}$ into spatial and temporal main effects and an interaction, with all terms being modelled by GMRF priors. More recently, Rushworth et al. (2014) utilise the autoregressive decomposition

$$f(\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_T) = f(\boldsymbol{\phi}_1) \prod_{j=2}^{T} f(\boldsymbol{\phi}_j | \boldsymbol{\phi}_{j-1}), \tag{3}$$

where $\boldsymbol{\phi}_j = (\phi_{1j}, \ldots, \phi_{Nj})$. They combine the decomposition (3) with the Leroux CAR prior for each $\boldsymbol{\phi}_j$, so that $\boldsymbol{\phi}_1$ is modelled by (2), and $\boldsymbol{\phi}_j \sim \mathrm{N}\left(\alpha\boldsymbol{\phi}_{j-1},\ \tau^2\mathbf{Q}(\rho, \mathbf{W})^{-1}\right)$ for $j = 2, \ldots, T$, where $\alpha \in [0,1]$ controls temporal autocorrelation. The global nature of the spatial autocorrelation induced by (2) for purely spatial random effects $(\phi_1, \ldots, \phi_N)$ can be seen from their theoretical partial autocorrelations:

$$\mathrm{Corr}[\phi_i, \phi_k | \boldsymbol{\phi}_{-ik}, \rho, \mathbf{W}] = \frac{\rho w_{ik}}{\sqrt{(\rho \sum_{r=1}^{N} w_{ir} + 1 - \rho)(\rho \sum_{s=1}^{N} w_{ks} + 1 - \rho)}}. \tag{4}$$

Under models (2) and (3), random effects for all pairs of neighbouring areal units (for which $w_{ik} = 1$) will be partially autocorrelated, and the strength of that partial autocorrelation will be controlled by $\rho$. Thus as $\rho$ will often be close to one (the spatial residual surfaces are autocorrelated as described in Section 2), a pair of adjacent areas exhibiting substantially different levels of unexplained risk will have those risks wrongly smoothed towards each other, masking the step change to be identified.

This prompted the development of spatially adaptive smoothing models for spatial data, and the extension to the spatio-temporal domain is one of the contributions of this paper. Brewer and Nolan (2007) and Reich and Hodges (2008) extend GMRF models by allowing the variance $\tau^2$ to vary across the study region, while Lawson et al. (2012), Charras-Garrido et al. (2012), Wakefield and Kim (2013) and Anderson et al. (2014) utilise a piecewise constant cluster model in the linear predictor to model step changes between neighbouring areas. Alternatively, Lu et al. (2007), Brezger et al. (2007), Ma et al. (2010), Lee and Mitchell (2013) and Lee et al. (2014) generalise CAR models by treating the non-zero elements of the adjacency matrix $\mathbf{W}$ as random variables. Under this approach, equation (4) implies that spatially adjacent $(\phi_i, \phi_k)$ can be partially autocorrelated or conditionally independent, depending on the estimated value of $w_{ik}$.

## 4. Methodology

### *4.1. General approach*

We present a novel spatially adaptive smoothing model for spatio-temporal data, which allows step changes to occur between adjacent areal units in the unexplained component of risk while treating their locations as unknown. This is achieved by modelling the adjacency elements in $\mathbf{W}$, that is $\mathbf{w}^+ = \{w_{ik}|i \sim k\}$ (of length $N_W = \mathbf{1}^T W \mathbf{1}/2$), as random variables on the unit interval, rather than being fixed equal to one. The remaining elements of $\mathbf{W}$ corresponding to non-adjacent areal units remain fixed at zero. Equation (4) shows that when $\rho$ is close to 1, then estimating $w_{ik} \in \mathbf{w}^+$ as close to 1 results in partial autocorrelation and hence smoothing between the spatially adjacent $(\phi_{ij}, \phi_{kj})$ for all time periods $j$. Conversely, if $w_{ik}$ is estimated as close to zero then $(\phi_{ij}, \phi_{kj})$ are close to conditionally independent for all time periods $j$, and no such spatial smoothing is enforced. In the latter case, a step change is said to exist in the random effects surface between areal units $(i, k)$ for all time periods $j$. We follow Lu and Carlin (2005) and quantify the evidence for a step change using

$$p_{ik} = \mathbb{P}(w_{ik} < 0.5|\mathbf{Y}), \tag{5}$$

the posterior probability of $w_{ik}$ being less than 0.5. Our proposed model is one of the first adaptive (localised) smoothing models for spatio-temporal data, and is outlined in two stages below.

### *4.2. Level 1 - Likelihood and random effects model for* $(Y_{ij}, \phi_{ij})$

The first level of our proposed model is given by

$$\begin{aligned}
Y_{ij}|E_{ij}, R_{ij} &\sim \text{Poisson}(E_{ij}R_{ij}) \quad i = 1, \ldots, N, \quad j = 1, \ldots, T, \\
\ln(R_{ij}) &= \mathbf{x}_{ij}^{\top}\boldsymbol{\beta} + \phi_{ij}, \\
\beta_0 &\sim \text{N}(0, 10000), \\
\boldsymbol{\phi}_1 &\sim \text{N}\left(\mathbf{0},\ \tau^2\mathbf{Q}(\mathbf{W}, \epsilon)^{-1}\right), \\
\boldsymbol{\phi}_j|\boldsymbol{\phi}_{j-1} &\sim \text{N}\left(\alpha\boldsymbol{\phi}_{j-1},\ \tau^2\mathbf{Q}(\mathbf{W}, \epsilon)^{-1}\right) \text{ for } j = 2, \ldots, T, \\
\tau^2 &\sim \text{Inverse-Gamma}(0.001, 0.001), \\
\alpha &\sim \text{Uniform}(0, 1).
\end{aligned} \tag{6}$$

The only difference from the model proposed by Rushworth et al. (2014) is that the GMRF prior proposed by Leroux et al. (1999) is replaced by the intrinsic GMRF prior (where $\rho = 1$), which is enforced because attempting to estimate $\rho$ and $\mathbf{W}$ could result in high posterior correlation and multimodality, because the random effects are spatially independent if either $\rho = 0$ or all elements of $\mathbf{w}^{+}$ equal zero. To avoid rank-deficiency of the precision matrix and subsequent problems with matrix inversion, the adjusted specification $\mathbf{Q}(\mathbf{W}, \epsilon) = \text{diag}(\mathbf{W1}) - \mathbf{W} + \epsilon\mathbf{I}$ is used, where $\epsilon\mathbf{I}$ is added to ensure that $\mathbf{Q}(\mathbf{W}, \epsilon)$ is diagonally dominant and hence invertible. This invertibility condition is required because in the second level of the model described below, elements in $\mathbf{W}$ are treated as random variables, necessitating the evaluation of the normalised prior density $f(\boldsymbol{\phi}_j|\boldsymbol{\phi}_{j-1})$. Sensitivity to the value of $\epsilon$ was checked in an initial modelling step, and was found not to affect estimation until $\epsilon$ was

increased to a relatively large value, such as $\epsilon > 10^{-2}$. Therefore in this paper we set $\epsilon = 10^{-7}$.

### *4.3. Level 2 - Adjacency model for elements in* $\mathbf{w}^+$

Our methodological contribution extends the model of Rushworth et al. (2014) by treating the elements $\mathbf{w}^+$ as binary random quantities on the unit interval, rather than being fixed at one. Specifying a continuous domain for $\mathbf{w}^+$ allows the direct application of a second stage GMRF prior, avoiding the need for a discrete prior such as the Ising model, for which no polynomial time algorithm exists to compute its normalising constant. We model $\mathbf{w}^+$ on the logit scale, $\mathbf{v}^+ = \log\left(\mathbf{w}^+/(\mathbf{1} - \mathbf{w}^+)\right)$, which has the back transformation $\mathbf{w}^+ = \exp(\mathbf{v}^+)/(1 + \exp(\mathbf{v}^+))$. The GMRF prior we propose for $\mathbf{v}^+$ has a constant mean $\mu$, a constant variance $\zeta^2$, and a precision matrix defined by the GMRF prior proposed by Leroux et al. (1999). This second stage GMRF prior requires us to specify an adjacency structure for the elements in $\mathbf{v}^+$, and here $v_{ik}, v_{rs} \in \mathbf{v}^+$ are defined as adjacent (denoted $ik \sim rs$) if the geographical borders they represent in the study region share a common vertex. Using this notation, we propose the following GMRF prior for $\mathbf{v}^+$:

$$\begin{aligned} p(\mathbf{v}^+|\zeta^2, \rho, \mu) &\propto \exp\left[-\frac{1}{2\zeta^2}\left(\rho \sum_{ik\sim rs}(v_{ik} - v_{rs})^2 + (1-\rho)\sum_{v_{ik}\in\mathbf{v}^+}(v_{ik}-\mu)^2\right)\right] \quad (7) \\ \zeta^2 &\sim \text{Inverse-Gamma}(0.001, 0.001), \\ \rho &\sim \text{Uniform}(0,1). \end{aligned}$$

This form highlights the role of $\rho$, which controls the extent to which step changes appear spatially clustered together around common vertices. When $\rho \approx 1$ the random variable $v_{ik}$, which controls the existence of a step change between areal units $(i, k)$, is smoothed spatially towards adjacent $v_{rs}$ via the penalty $\sum_{ik\sim rs}(v_{ik} - v_{rs})^2$, which thus induces spatially clustered step changes, this model is conceptually similar to the 'CAR2' model proposed in Ma et al. (2010). In contrast, when $\rho \approx 0$ each $v_{ik}$ is smoothed non-spatially towards the overall mean $\mu$ by the penalty $\sum_{v_{ik}\in\mathbf{v}^+}(v_{ik}-\mu)^2$,

which does not encourage spatial clustering of step changes. In order to avoid numerical problems when transforming between $\mathbf{v}^+$ and $\mathbf{w}^+$, the sample space for each $v_{ik} \in \mathbf{v}^+$ is truncated to the interval $[-15, 15]$, yielding a sample space of $[0.000000306, 0.9999997]$ for $w_{ik}$, which is close to the intended $[0, 1]$ interval.

The prior mean $\mu$ is fixed in (7), to ensure that the induced prior on the untransformed $\mathbf{w}^+$ scale is consistent with our prior beliefs about the prevalence of step changes. Specifically, given the level of spatial autocorrelation evident in the residuals shown in Figure 1, and the associated Moran I statistics reported in Section 2, one would expect there to be relatively few step changes in the random effects surface. In order to be consistent with this preference we choose $\mu > 0$, as choosing $\mu < 0$ implies a marginal mean for $w_{ik}$ of less than $\exp(0)/(1 + \exp(0)) = 0.5$. However, Figure 2 shows that the induced prior distribution for $w_{ik}$ depends on $\zeta^2$ as well as $\mu$, with the left and right panels showing $\mu = 0$ and $\mu = 15$ respectively for various values of $\zeta$. When $\mu = 0$ the prior density for $w_{ik}$ can have a mode at 0.5, which is incongruous with our prior beliefs about $w_{ik}$ being close to one for most $w_{ik} \in \mathbf{w}^+$. Initial simulations confirmed that setting $\mu = 0$ leads to spurious step changes being identified. In contrast, when $\mu = 15$ the prior assigns high prior probability to $w_{ik} \approx 0$ or $w_{ik} \approx 1$ or both, with little prior probability in between. The ratio of the densities at $\{0, 1\}$ depends on $\zeta$, so that when $\zeta$ is small, almost all prior mass is concentrated around $w_{ik} = 1$, hence strongly discouraging boundaries. In contrast, as $\zeta$ increases the prior becomes more symmetric and 'U' shaped, with equal point masses at 0 and 1 expressing ambivalence about the presence or absence of step changes. Thus fixing $\mu = 15$ ensures that clear step change decisions, that is $w_{ik}$ close to zero or one, are preferred over ambiguous values such as $w_{ik} = 0.5$.

### *4.4. Inference*

Our model is fitted within a Bayesian framework using Markov-chain Monte Carlo (McMC) sampling, and is freely available via the `R` package `CARBayesST` (Lee et al. (2015)). The McMC algorithm is a combination of Gibbs and Metropolis-Hastings steps, and the high dimensionality of the parameter space requires an algorithm

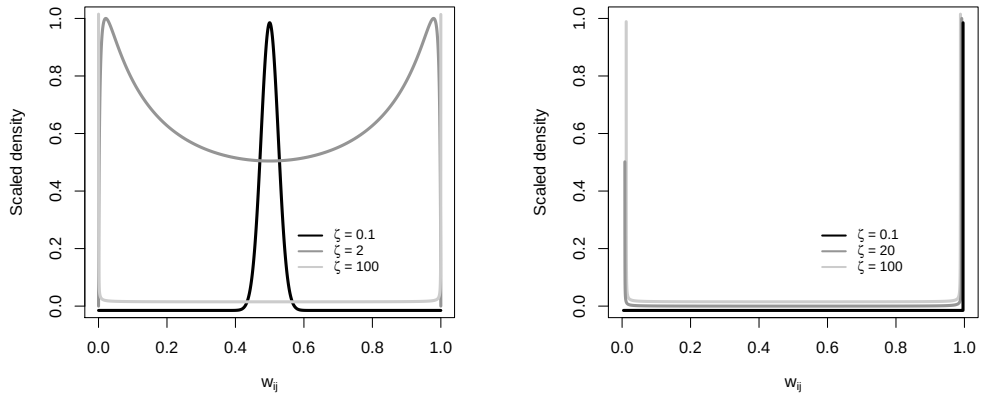


**Fig. 2.** Plots showing scaled prior densities for $w_{ij}$ for prior means $\mu = 0$ (left) and $\mu = 15$ (right). In each plot the densities resulting from different $\zeta$ values are shown by different coloured lines.

that minimises the overall computational burden. We exploit matrix sparsity resulting from the GMRF prior precision matrix wherever possible, using efficient loops written in C++ and matrix triplet form to perform the algebraic manipulations. The random effects were updated using fast one-at-a-time Metropolis-Hastings steps rather than using a joint or block updating scheme. While joint updates (for example Knorr-Held and Rue (2002)) can achieve substantially better mixing for an equivalent number of McMC steps and particularly when the parameters are highly correlated, the computational cost was found to outweigh the benefits of better sampling efficiency. Further details are given in the supporting information accompanying this paper.

## 5. Simulation study

In this section we comprehensively test the performance of two variants of the proposed model on simulated data under a range of scenarios, and compare their performance against two commonly used alternatives. The two existing models are those proposed by Knorr-Held (2000) (denoted **(1)**) and Rushworth et al. (2014) (denoted **(2)**), although the former is implemented with GMRF priors proposed by Leroux

et al. (1999) rather than a convolution of independent and intrinsic GMRF priors. Additionally, model **(1)** is implemented with independent interactions (Gaussian with zero-mean and a common variance), although 3 others types of interactions were proposed by Knorr-Held (2000). For more details see the vignette accompanying the `CARBayesST` software. Model **(3)** is the adaptive smoothing model proposed here with the simplification that $\rho = 0$, while Model **(4)** is the full model where $\rho$ is not treated as fixed. Model **(3)** *a-priori* treats each $w_{ik} \in \mathbf{w}^+$ independently, and therefore does not encourage configurations in which step changes cluster together. Our primary focus in this study is to assess the ability of each model to estimate the spatio-temporal pattern in disease risk, and identify step changes in risk between neighbouring areas.

### *5.1. Data generation and study design*

Simulated disease counts $\{Y_{ij}\}$ are generated for the England study region from a Poisson log-linear model, that is $Y_{ij} \sim \text{Poisson}(E_{ij}R_{ij})$, where both the size of the expected counts and the number of time periods $T$ are varied to assess their impact on model performance. The log-risk surfaces are generated for each time period from a multivariate Gaussian distribution, whose precision matrix is defined by the intrinsic GMRF prior (Besag, 1974, Besag et al., 1991) and hence produces spatially smooth surfaces. To simulate spatial step changes in log-risk, a piecewise constant mean surface is specified for the random effects, which is displayed in the left panel of Figure 3. Lighter shaded areas exhibit a mean risk of 1 while the darker shaded areas have a mean risk level of $A$, and the black lines correspond to the locations of true step changes. An example realisation of this surface is shown in the right panel of Figure 3 for $A = 1.5$, where the clusters of high-risk areas are evident. To ensure the true risk surface is not identical for all time periods, independent random noise is added to the risk in each areal unit for each time period. The scenarios considered in this study are summarised in Table 1, which shows that we consider $T = 1, 5, 20$ time periods, elevated risk levels of $A = 1, 1.5, 2$, and disease prevalences of $E = 10, 50, 100$. For the $A = 1$ scenario this corresponds to a spatially smooth risk surface with no step changes, which tests the model's propensity for identifying

**Table 1.** Description of the scenarios in the simulation study.

| Scenario type | Parameters varied | Parameters fixed |
|---|---|---|
| Varying time dimension | $T \in \{1, 5, 20\}$ | $a = 1.5$; $E = 75$ |
| Relative risk in high regions | $A \in \{1, 1.5, 2\}$ | $T = 5$; $E = 75$ |
| Expected cases | $E \in \{10, 50, 100\}$ | $T = 5$; $a = 1.5$ |

step changes when none are present (false positives).

**Fig. 3.** Left: Locations of the true step changes in risk, illustrated by black lines following the borders between the selected subregions. Darker shading indicates areas with true risk of 1.5 while lighter shading indicates a true risk of 1. Right: A single realisation of the spatial risk surface assuming $A = 1.5$.

### *5.2. Results*

One hundred data sets were generated under each of the 9 scenarios shown in Table 1, and Models **(1)** - **(4)** were fitted in turn. Inference for each model was based on 30,000 McMC samples following a burn-in period of 20,000 samples, after which convergence was assessed to have been reached. The quality of the estimation of the spatio-temporal pattern in disease risk was quantified by its root-mean squared error, RMSE= $\sqrt{\frac{1}{NT}\sum_{i,j}(R_{ij} - \hat{R}_{ij})^2}$, as well as by the Deviance Information Criterion (DIC), effective number of parameters ($pD$) and the coverage probabilities of

the 95% credible intervals. Receiver-operating characteristic (ROC) curves were also computed to quantify the accuracy of the step change detection, which were based on each models sensitivity and specificity at identifying true step changes. These statistics compared $\mathbb{E}[w_{ij}|\mathbf{Y}]$ to a threshold value $p^*$, where if $\mathbb{E}[w_{ij}|\mathbf{Y}] < p^*$ a step change was identified where as for the converse no step change was declared. The value of $p^*$ was varied from 0 to 1 at intervals of 0.01, and the ROC curve is a plot of sensitivity against specificity. However, for ease of presentation the Area Under the Curve (AUC) is presented here rather than the full ROC curve, and an AUC of one corresponds to perfect step change identification. AUC associated with step-change estimation was only available for Models **(3)** and **(4)**, as Models **(1)** and **(2)** do not estimate step-changes as $\mathbf{w}^+$ is fixed and not estimated in these models. Therefore, AUC is of interest to validate and compare the performance the adaptive models proposed.

Table 2 shows the RMSE, DIC, $pD$ and coverages associated with each model across the nine simulation scenarios, from which a number of patterns emerge. Firstly, the simplified adaptive model **(3)** with no spatial smoothing across step changes generally outperforms the non-adaptive models **(1)** and **(2)** and the full adaptive model **(4)** in terms of lower RMSE, DIC and $pD$. In the $E = 10$ scenario model **(1)** clearly outperforms **(3)** and **(4)** in terms of RMSE, DIC and coverage which is likely the result of incorrect boundary identification from the adaptive models (see 3). Also in the $A = 1$ scenario with no true step changes model **(4)** exhibits a very low $pD$, which indicates oversmoothing caused by false identification of step changes in the surface (see 3). Overall, Table 2 strongly suggests that the spatial smoothing imposed on $\mathbf{w}^+$ by **(4)** is sub-optimal compared to assuming each element $w_{ik} \in \mathbf{w}^+$ is *a-priori* independent as in **(3)**. For all models, RMSE decreases as both the number of time periods $T$ and disease prevalence $E$ increases, which is due to an increase in the amount of data. Confidence interval coverage was generally very good, with all coverage levels varying between 93% and 97%, excepting $A = 1$ when the cover-

**Table 2.** Median root mean-squared error (RMSE), 95% credible interval coverages associated with the fitted risks, Deviance Information Criterion (DIC), and the effective number of parameters (pD) for each model and scenario.

| | | **RMSE** | | | | **DIC** | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | **(1)** | **(2)** | **(3)** | **(4)** | **(1)** | **(2)** | **(3)** | **(4)** |
| **Time** | $T = 1$ | 0.092 | 0.092 | 0.088 | 0.090 | 2497 | 2497 | 2489 | 2502 |
| | $T = 5$ | 0.052 | 0.070 | 0.034 | 0.062 | 11986 | 12116 | 11815 | 12132 |
| | $T = 20$ | 0.036 | 0.058 | 0.032 | 0.036 | 47355 | 47690 | 47217 | 47325 |
| **Risk** | $A = 1$ | 0.023 | 0.027 | 0.024 | 0.024 | 11598 | 11625 | 11608 | 11614 |
| | $A = 1.5$ | 0.052 | 0.071 | 0.035 | 0.062 | 11983 | 12116 | 11815 | 12107 |
| | $A = 2$ | 0.059 | 0.089 | 0.039 | 0.059 | 12131 | 12442 | 11923 | 12200 |
| **Cases** | $E = 10$ | 0.100 | 0.129 | 0.111 | 0.114 | 8549 | 8570 | 8577 | 8595 |
| | $E = 50$ | 0.061 | 0.079 | 0.043 | 0.069 | 11293 | 11393 | 11164 | 11407 |
| | $E = 100$ | 0.046 | 0.063 | 0.031 | 0.051 | 12472 | 12634 | 12284 | 12575 |

| | | **pD** | | | | **Coverage** | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | **(1)** | **(2)** | **(3)** | **(4)** | **(1)** | **(2)** | **(3)** | **(4)** |
| **Time** | $T = 1$ | 166 | 166 | 155 | 164 | 0.930 | 0.930 | 0.955 | 0.950 |
| | $T = 5$ | 322 | 526 | 159 | 520 | 0.970 | 0.940 | 0.980 | 0.960 |
| | $T = 20$ | 565 | 1322 | 351 | 377 | 0.960 | 0.910 | 0.910 | 0.850 |
| **Risk** | $A = 1$ | 98 | 234 | 36 | 5 | 0.986 | 0.996 | 0.878 | 0.263 |
| | $A = 1.5$ | 323 | 528 | 164 | 475 | 0.963 | 0.930 | 0.979 | 0.956 |
| | $A = 2$ | 363 | 761 | 173 | 475 | 0.963 | 0.941 | 0.982 | 0.964 |
| **Cases** | $E = 10$ | 159 | 274 | 186 | 201 | 0.964 | 0.941 | 0.950 | 0.944 |
| | $E = 50$ | 290 | 456 | 154 | 413 | 0.962 | 0.932 | 0.985 | 0.963 |
| | $E = 100$ | 347 | 584 | 171 | 480 | 0.965 | 0.936 | 0.978 | 0.956 |

**Table 3.** Median area under the ROC curve for step change identification across 100 simulations for models **(3)** and **(4)**. Bracketed figures correspond to the 10% quantile of areas. For $A = 1$ SPF denotes the specificity since there are no true step changes to identify in this scenario.

| | | **Median area under ROC curves** | |
|---|---|---|---|
| | | **(3)** | **(4)** |
| **Temporal replication** | $T = 1$ | 0.7399 (0.6696) | 0.4995 (0.4898) |
| | $T = 5$ | 0.9999 (0.9996) | 0.4995 (0.4988) |
| | $T = 20$ | 0.9997 (0.9988) | 0.9151 (0.4995) |
| **Relative risk** | $A = 1$, SPF | 0.9769 (0.9463) | 0.4186 (0.2649) |
| | $A = 1.5$ | 0.9993 (0.9996) | 0.5913 (0.4988) |
| | $A = 2$ | 0.9999 (0.9997) | 0.7746 (0.4995) |
| **Expected cases** | $E = 10$ | 0.6672 (0.6262) | 0.4979 (0.4885) |
| | $E = 50$ | 0.9925 (0.9780) | 0.5801 (0.4988) |
| | $E = 100$ | 0.9999 (0.9998) | 0.6463 (0.4994) |

ages were 0.986, 0.996, 0.878 and 0.263 for models **(1)**, **(2)**, **(3)** and **(4)**, respectively.

Table 3 displays the median AUC statistic across the set of ROC curves calculated for each scenario. The numbers in brackets are the tenth percentile of that distribution, and summarise the variation across the 100 simulated data sets. An exception to this is the $A = 1$ scenario, which displays the specificity because as the risk surface is spatially smooth there are no true step changes to identify. For model **(3)** the median AUC values are close to the maximal value of 1, indicating very accurate step change identification. The exception to this occurs when $T = 1$ (AUC = 0.7399) and $E = 10$ (AUC = 0.6672), which results from limited information about step change location provided by the data in both cases. In contrast, Model **(4)** always performs much less well, with median AUC values ranging between 0.5 and 0.9151. The poor performance of Model **(4)** is also evident in the tenth percentile results, and re-enforces the RMSE and coverage results displayed in Table 2. It is likely to be because model **(4)** forces the step changes to be spatially smooth, thus inducing a set of false positives that are spatially close to the real boundaries. The other main result from Table 3 is that the AUC increases as the number of time periods $T$ increases and as the size of the expected cases increases, which is due to an increase in the amount of data.

## 6. Results of the England case study

The simulation study has shown that the simplified adaptive model with $\rho = 0$ (Model **(3)**) is capable of identifying spatial discontinuities where they exist, where as the full model (Model **(4)**) is not, Additionally Model **(3)** fits the data better than Model **(4)** in terms of DIC, and thus the latter is not considered here. We thus apply three models to the England circulatory and respiratory data sets, the interaction model of Knorr-Held (2000) (Model **(1)**), the global smoothing model of Rushworth et al. (2014) (Model **(2)**) and the adaptive smoothing model proposed here with the simplification that $\rho = 0$ (Model **(3)**). The three covariates discussed in Section 2 are included in each model, which are the proportion of working age people claiming Job Seekers Allowance (JSA), the proportion of each LUA classified as urban (Urbanity) and the particulate matter concentrations ($PM_{10}$). Inference for all models is based on thinning (by 10) $10^6$ posterior samples including a burn-in period of a further $10^5$ samples. In analysing these data our goals are to: (i) estimate the spatio-temporal pattern in disease risk to quantify the extent of health inequalities; and (ii) estimate the location of any step changes in the unexplained spatial risk structure, which will assist in the identification of unmeasured confounders. The analysis provided here can be reproduced using the code and data in the supporting material accompanying this paper.

### *6.1. Model fit*

The top two rows of Table 4 display the overall fit of each model to each data set, by presenting the DIC and the effective number of parameters $pD$. It shows that the adaptive smoothing model **(3)** fits the data better than the global smoothing models **(1)**, **(2)** for both diseases, with reductions in DIC in both cases. Additionally, model **(3)** has a markedly smaller number of effective parameters ($pD$) than models **(1)**, **(2)**, despite having a more complex specification. This is because its ability to identify step changes permits the GMRF component to smooth more strongly elsewhere in the spatial surface, resulting in smaller variance estimates for $\tau^2$ than from model **(2)**. This implies a greater level of penalisation of the random effects

**Table 4.** Diagnostics for models **(1)** to **(3)** for the England circulatory and respiratory admissions data sets.

| | Circulatory | | | Respiratory | | |
|---|---|---|---|---|---|---|
| **Diagnostic** | **(1)** | **(2)** | **(3)** | **(1)** | **(2)** | **(3)** |
| DIC | 35,269 | 35,314 | 35,056 | 35,223 | 35,239 | 34,993 |
| $pD$ | 2,783 | 2,890 | 2630 | 2,825 | 2,887 | 2,629 |
| % of borders with $p_{ik} > 0.75$ | - | - | 30.8 | - | - | 32.6 |
| % of borders with $p_{ik} > 0.99$ | - | - | 14.9 | - | - | 17.1 |
| $\hat{\tau}^2$ | - | 0.0295 | 0.0127 | - | 0.0395 | 0.0130 |
| $\hat{\zeta}^2$ | - | - | 250.4 | - | - | 254.9 |
| $\hat{\alpha}$ | - | 0.963 | 0.961 | - | 0.969 | 0.964 |

and hence a reduction in the overall $pD$.

### *6.2. Covariate effects*

The effects of the three covariates on circulatory and respiratory disease risks estimated by model **(3)**, the best fitting model, are displayed in Table 5, where all results are presented as relative risks associated with one standard deviation increases in each covariate. Standard deviations were calculated using the raw covariate data, and are shown in the second column of Table 5. The table shows that increasing socio-economic deprivation (as measured by JSA) increases risk for hospital admission of both diseases, with relative risks of 1.08 (circulatory) and 1.135 (respiratory) respectively. This harmful effect of poverty is well known, and is a proxy for differences in average lifestyle such as prevalence of smoking, drinking, exercise etc. Both diseases also show substantial effects of urbanicity on disease risk, with more urban areas exhibiting increased risks of 1.096 (circulatory) and 1.142 (respiratory) respectively. Finally, particulate matter air pollution appears to have no impact on circulatory disease risk and a slight effect on respiratory disease risk, with the latter relative risk being 1.021.

### *6.3. Health inequalities*

Maps of the average risks across all years from model **(3)** are displayed in the left column of Figure 4, and show similar spatial patterns in risk for both diseases, with a Pearson's correlation coefficient of 0.892. The maps show that the average risk varies over space with values between (0.433, 1.636) and (0.175, 2.147) respectively

**Table 5.** Relative risks and associated 95% credible intervals associated with 1 standard deviation increases in each covariate, as estimated by model **(3)**

| | | Circulatory | |
|---|---|---|---|
| **Covariate** | **St.Dev** | **RR** | **95% CI** |
| **JSA** | 0.0476 | 1.0799 | (1.0586, 1.1021) |
| **Urbanicity** | 0.257 | 1.0964 | (1.0757, 1.1173) |
| **$PM_{10}$** | 2.903 | 0.9954 | (0.9730, 1.0209) |
| | | **Respiratory** | |
| **Covariate** | **St.Dev** | **RR** | **95% CI** |
| **JSA** | 0.0476 | 1.1346 | (1.1126, 1.1555) |
| **Urbanicity** | 0.257 | 1.1418 | (1.1217, 1.1649) |
| **$PM_{10}$** | 2.903 | 1.0211 | (0.9946, 1.0557) |

for circulatory and respiratory disease, suggesting the presence of substantial health inequalities. These inequalities have generally widened over time, as the difference between the highest and lowest respiratory disease risk was 1.77 in 2001 and 2.13 in 2010. For circulatory disease a similar pattern is evident, with an estimated difference between highest and lowest risk of 1.39 in 2001 and 1.54 in 2010.

### *6.4. Step change identification*

Table 4 summarises the number of step changes in the unexplained component of the risk surface, based on $p_{ik} = \mathbb{P}(w_{ik} < 0.5|\mathbf{Y})$ values above a threshold of 0.75 and 0.99. The higher 0.99 level threshold was used by Lu and Carlin (2005), and results in 15.1% of borders being step changes for circulatory disease and 17.1% for respiratory disease. These step changes are largely similar between the diseases, with 92% agreement between their locations. They are displayed in the right column of Figure 4 as white lines, while the grey shading represents the time averaged exponentiated random effects surface which corresponds to the unexplained component of the variation in disease risk. The figure shows evidence of much higher unexplained risks of hospital admission in areal units containing large cities, and in the central band of Northern England that incorporates Manchester and Yorkshire, even after adjusting for the covariates. It is striking that these features are largely consistent between the two diseases, so that although the estimated risks have different overall magnitudes, they exhibit very similar spatial patterns. Public health professionals can use these results to identify potential risk factors for disease, by searching for risk factors that

exhibit step changes in the same locations as those exhibited in Figure 4.

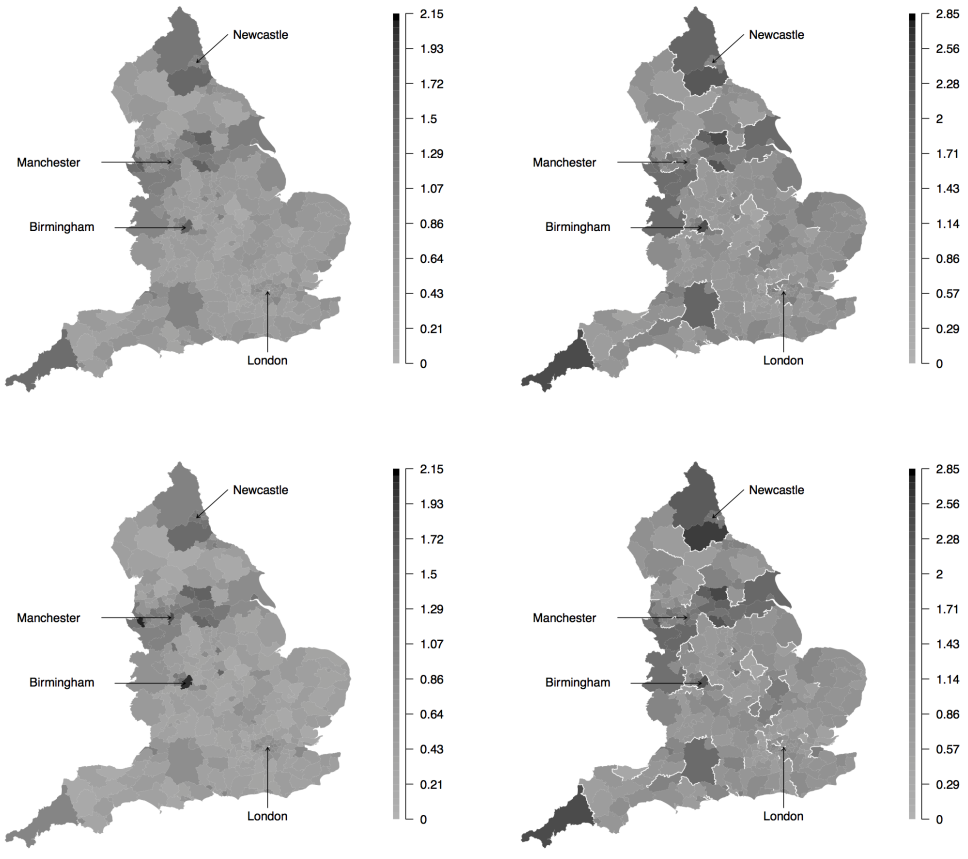


**Fig. 4.** Maps showing the average risk surface (left column) and the unexplained component of the risk surface (right column) for both diseases from model **(3)**. The top row relates to circulatory disease while the bottom row relates to respiratory disease. The white lines on the maps in the right column correspond to step changes that have been identified using a cutoff of $p_{ik} \geq 0.99$ in (5).

### *6.5. Sensitivity analysis*

A prior sensitivity analysis was performed to test the robustness of the above results to changes in prior specifications. In particular, the shape and scale parameters in the hyperpriors for $\zeta^2$ and $\tau^2$ were changed to (1,1) and (1, 0.001), and the prior variance for $\boldsymbol{\beta}$ was reduced to 100. The results reported above were found to be robust to these modifications. Our choice of Inverse-Gamma priors here is

due to their conjugacy, but alternatives could be explored such as the half-Cauchy distribution as suggested by Gelman et al. (2006).

## 7. Discussion

In this paper a new study of the spatio-temporal structure of circulatory and respiratory disease risk in England was presented, with the goal of understanding the extent of health inequalities and whether the data present evidence of step changes in disease risk between pairs of adjacent regions. Consequently, a new spatially adaptive smoothing model was developed, that can estimate the location and strength of such step-changes. The model is a spatially adaptive extension to the class of GMRF prior distributions, and is one of the first models for step change identification in spatio-temporal disease risk. Freely available software via the `CARBayesST` package for `R` is provided to allow others to utilise our model, and this is one of the first `R` packages for spatio-temporal areal unit modelling.

The simulation study in Section 5 established the superiority of our model over global smoothing alternatives, in terms of both risk estimation and model fit. Our model was successful at recovering the locations of known step changes in simulated data, with AUC statistics close to one for a range of different scenarios. These AUC statistics were higher if the step changes were assumed to be independent in space, because *a-priori* assuming spatial clustering resulted in false step changes being identified close to real step changes. This is an interesting result, as one may have naively assumed that for spatial data the step changes should be modelled spatially, which we have shown is not the case. Thus existing global smoothing models are sub-optimal for space time disease mapping in two respects, they smooth over such step changes leading to poorer estimation of disease risk, and they cannot identify such step changes which themselves provides etiological evidence about potential unmeasured risk factors.

Section 6 provided strong evidence of step changes in the unexplained component

of risk for the England data. Additionally, better model fit with a smaller number of effective parameters was observed compared to the global smoothing models, due to increased levels of smoothing in locations where step changes were absent. Thus non-adaptive smoothing models may overfit some data sets, by imposing too weak a spatial smoothing constraint due to the presence of step changes in risk. A striking association was found between the fitted risks and identified step changes between circulatory and respiratory disease, perhaps indicating the influence of the same unobserved risk factors. Therefore in future work we will try and identify such unmeasured confounders, to see if they are indeed common to both diseases. A univariate approach to modelling was taken here as discussed in Section 2, but the between disease correlation naturally suggests a bivariate spatio-temporal model as a future avenue of work. Such models have yet to be proposed in a spatio-temporal context, and would require the extension of multivariate spatial models such as multivariate conditional autoregressive (Gelfand and Vounatsou, 2003) models and shared component models (Knorr-Held and Best, 2001). A final avenue for future work is to use the model in an ecological regression context, where the effect of an exposure on disease risk is of primary interest rather than the spatio-temporal pattern in disease risk. The efficacy of adaptive smoothing models in this context may be to reduce spatial confounding between the random effects and the covariates as suggested by Clayton et al. (1993), due to the reduction in the random effects variance (see Table 4), and environmental factors such as air pollution would be a natural context for such work.

## 8. Acknowledgements

This work was funded by the Engineering and Physical Sciences Research Council (EPSRC), via grant EP/J017442/1, and the authors would like to thank the associate editor and a reviewer for their constructive comments that greatly improved the paper.

## References

Anderson, C., D. Lee, and N. Dean (2014). Identifying clusters in Bayesian disease mapping. *Biostatistics 15*, 457–469.

Bernardinelli, L., D. Clayton, C. Pascutto, C. Montomoli, M. Ghislandi, and M. Songini (1995). Bayesian analysis of spacetime variation in disease risk. *Statistics in Medicine 14*(21-22), 2433–2443.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 192–236.

Besag, J., J. York, and A. Mollié (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics 43*(1), 1–20.

Brewer, M. J. and A. J. Nolan (2007). Variable smoothing in Bayesian intrinsic autoregressions. *Environmetrics 18*(8), 841–857.

Brezger, A., L. Fahrmeir, and A. Hennerfeind (2007). Adaptive Gaussian Markov random fields with applications in human brain mapping. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 56*(3), 327–345.

Charras-Garrido, M., D. Abrial, J. De Goër, S. Dachian, and N. Peyrard (2012). Classification method for disease risk mapping based on discrete hidden Markov random fields. *Biostatistics 13*(2), 241–255.

Clayton, D., L. Bernardinelli, and C. Montomoli (1993). Spatial Correlation in Ecological Analysis. *Int J Epidemiol 22*, 1193–1202, DOI:10.1093/ije/22.6.1193.

Gelfand, A. E. and P. Vounatsou (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics 4*(1), 11–15.

Gelman, A. et al. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian analysis 1*(3), 515–534.

Green, P. J. and S. Richardson (2002). Hidden Markov models and disease mapping. *Journal of the American statistical association 97*(460), 1055–1070.

Knorr-Held, L. (2000). Bayesian modelling of inseparable space-time variation in disease risk. *Statistics in Medicine 19*(17-18), 2555–2567.

Knorr-Held, L. and N. G. Best (2001). A shared component model for detecting joint and selective clustering of two diseases. *Journal of the Royal Statistical Society. Series A (Statistics in Society) 164*(1), pp. 73–85.

Knorr-Held, L. and H. Rue (2002). On block updating in markov random field models for disease mapping. *Scandinavian Journal of Statistics 29*(4), 597–614.

Lawson, A. B., J. Choi, B. Cai, M. Hossain, R. S. Kirby, and J. Liu (2012). Bayesian 2-stage space-time mixture modeling with spatial misalignment of the exposure in small area health data. *Journal of agricultural, biological, and environmental statistics 17*(3), 417–441.

Lee, D. and R. Mitchell (2013). Locally adaptive spatial smoothing using conditional auto-regressive models. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 62*(4), 593–608.

Lee, D. and R. Mitchell (2014). Controlling for localised spatio-temporal autocorrelation in long-term air pollution and health studies. *Statistical Methods in Medical Research 23*, 488–506.

Lee, D., A. Rushworth, and G. Napier (2015). *CARBayesST: Spatio-Temporal Generalised Linear Mixed Models for Areal Unit Data*. R package version 2.1.

Lee, D., A. Rushworth, and S. K. Sahu (2014). A Bayesian localized conditional autoregressive model for estimating the health effects of air pollution. *Biometrics 70*, 419–429.

Leroux, B. G., X. Lei, and N. Breslow (1999). *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, Chapter Estimation of Disease Rates in Small Areas: A new Mixed Model for Spatial Dependence, pp. 179–191. New York, NY: Springer New York.

Lu, H. and B. P. Carlin (2005). Bayesian areal wombling for geographical boundary analysis. *Geographical Analysis 37*(3), 265–285.

Lu, H., C. S. Reilly, S. Banerjee, and B. P. Carlin (2007). Bayesian areal wombling via adjacency modeling. *Environmental and Ecological Statistics 14*(4), 433–452.

Ma, H., B. P. Carlin, and S. Banerjee (2010). Hierarchical and joint site-edge methods for medicare hospice service region boundary analysis. *Biometrics 66*(2), 355–364.

MacNab, Y. C. and C. Dean (2001). Autoregressive spatial smoothing and temporal spline smoothing for mapping rates. *Biometrics 57*(3), 949–956.

Office for National Statistics (2014). Inequality in healthly life expectancy at birth by national deciles of area deprivation: England, 2009-11.

Reich, B. J. and J. S. Hodges (2008). Modeling longitudinal spatial periodontal data: A spatially adaptive model with tools for specifying priors and checking fit. *Biometrics 64*(3), 790–799.

Richardson, S., A. Thomson, N. Best, and P. Elliott (2004). Interpreting posterior relative risk rstimates in disease mappling studies. *Environmental Health Perspectives 112*, 1016–1025.

Rue, H. and L. Held (2005). *Gaussian Markov random fields: theory and applications*. CRC Press.

Rushworth, A., D. Lee, and R. Mitchell (2014). A spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in Greater London. *Spatial and Spatio-temporal Epidemiology 10*, 29–38.

Ugarte, M. D., T. Goicoa, and A. F. Militino (2010). Spatio-temporal modeling of mortality risks using penalized splines. *Environmetrics 21*(3-4), 270–289.

Wakefield, J. and A. Kim (2013). A Bayesian model for cluster detection. *Biostatistics 14*(4), 752–765.

Womble, W. (1951). Differential systematics. *Science 114*, 315–322.